Quantum Information and Computation, Vol. 0, No. 0 (2003) 000–000
© Rinton Press

arXiv:1510.00409v4 [physics.gen-ph] 9 May 2016

A linear time quantum algorithm for 3SAT

Zachary B. Walters[a]
University of California, Davis and Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley, California 94720, USA

Received (received date)
Revised (revised date)

A quantum algorithm for the NP complete problem of satisfying boolean formulas with three variables per clause (3SAT) is presented. Departing from traditional models of quantum computation, this algorithm makes extensive use of irreversible operations to incoherently transfer population from states which do not solve some problem of interest to states which do. Provided that a solution exists, the algorithm yields exponential decay of nonsolution probability, at a rate controlled by the user.

Keywords: quantum algorithm, 3SAT, NP complete

1 Introduction

An efficient method for solving NP complete problems, or the proof that no such method is possible, is a major open problem in theoretical computer science. Broadly speaking, computer science draws a distinction between problems which can be solved efficiently, using a number of steps which grows as a polynomial in the size of the input ("class P") and those which cannot. Within this basic distinction, a class of particular interest consists of problems whose answers can be checked for correctness but not necessarily found in polynomial time ("class NP"). It is believed, although not proven, that class NP is strictly larger than class P – that is, that there are some problems which can have no efficient solution algorithm, even though a candidate answer can be checked for correctness in polynomial time[1].

NP complete problems are problems within NP to which any other problem in NP can be reduced in polynomial time. Thus, a polynomial time algorithm for an NP complete problem would double as a polynomial time algorithm for any problem in NP. Conversely, if $P \neq NP$ as expected, no polynomial time algorithm using a classical computer can exist for the solution of any NP complete problem. Thousands of problems in science and mathematics have been shown to have this property [2, 3, 4, 5], making them computationally intractable for even the largest classical computing resources using known algorithms, and possibly intractable outright.

A quantum computer generalizes the behavior of its classical counterpart by operating, not on deterministic bits, but rather on mixtures of quantum states. In the classic Deutsch model [6], a quantum computation consists of a unitary operator applied to a initially pure quantum state. A small set of one- and two bit operators are universal, in the sense of being sufficient to emulate an arbitrary N bit unitary operation to arbitrary precision using a quantum mechanical analogue of a Turing machine[7, 8, 9]. The adiabatic model of computation

[a]email: ZBWalters@lbl.gov

[10] performs unitary operations by slowly changing the parameters of an N bit Hamiltonian in such a way as to project the initial wavefunction onto a wavefunction which solves the problem of interest. Both models are logically equivalent [11], with algorithmic running times related to one another by a polynomial factor.

The ability of a quantum system to occupy many states at once means that a quantum algorithm is trivially parallel, in the sense that all operations are performed upon every state represented in the system. However, the rules of quantum measurement limit the usefulness of this "quantum parallelism." Any measurement of the system's state is constrained to yield discrete values, and to destroy the amplitude for any state which does not correspond to the value measured. It is not possible to control the outcome of a measurement, except by preparing a wavepacket which is already an eigenstate of the measurement operator. Thus, any conditional branch in a quantum algorithm has the potential to destroy the amplitude corresponding to a desired solution state. Although many fewer quantum than classical algorithms have been discovered, it is known that quantum computation offers significant speedup over the fastest known classical algorithms for at least a few problems – not least of which, the simulation of other quantum systems [12, 13]. The best known of these are Shor's factoring algorithm [14, 15] and Grover's algorithm for searching an unordered database [16, 17]. A recent review of quantum algorithms is given in [18], and a current repository is maintained at [19].

To date no quantum algorithm has been discovered which can solve NP complete problems in polynomial time, and an important negative result [20] has shown that, relative to a uniformly chosen random oracle, a quantum Turing machine performing unitary transformations cannot solve arbitrary problems in NP in time $\mathbb{O}(2^{n/2})$. Limitations on the efficacy of adiabatic quantum algorithms are imposed by the existence of exponentially small energy gaps [21, 22, 23], requiring that the N bit Hamiltonian be changed over correspondingly large times in order to remain adiabatic.

This paper presents a linear time quantum algorithm for 3SAT, a canonical example of an NP complete problem. In contrast to the Deutsch and adiabatic models of quantum computation, this algorithm makes extensive use of irreversible quantum operations, so that it operates on a density matrix rather than a pure quantum state. Populations of nonsolution states are reduced by repeated decimation – if a given state fails to solve all or part of the problem, the states of the variables in question are rotated into their opposite value by a user defined rotation angle. Varying these rotation angles disrupts slowly decaying eigenvectors, so that the populations of states which fail to solve the problem decay exponentially to zero at a rate determined by the user. The problem of branching is treated by introducing temporary "scratch" bits, which are entangled with the states of particular bits, used to perform operations on others, and finally destroyed by nonselective measurement orthogonal to a chosen "axis of truth." This orthogonal measurement allows irreversible operations to be performed on the density matrix without extracting any information regarding the states of bits representing variables in the problem.

2 Background

The three variable Boolean satisfiability problem, or 3SAT, is the problem of determining whether a logical statement involving a set of variables can be satisfied by assigning each

variable to be true or false. The problem consists of a set of three variable clauses

$$(X_1 \vee X_2 \vee \neg X_3) \wedge (\neg X_2 \vee X_4 \vee X_5) \wedge \ldots \tag{1}$$

Within a clause, a variable may be represented by its value or by its negation, and variables within the same clause are joined by an OR operator. Different clauses are joined by an AND operator, and the same variable or its negation may appear in multiple clauses. The problem is solved if each variable can be assigned to be true or false in such a way that the statement as a whole evaluates to true.

In a quantum computer, the state of each variable is given by a superposition of two states

$$|X_n\rangle = \cos(\theta)\,|T\rangle + \sin(\theta)e^{i\phi}\,|F\rangle\,, \tag{2}$$

where the sum of the squared amplitudes yields 1. A density matrix has the form

$$\rho = \sum_{i,j} |\alpha_i\rangle\,\rho_{ij}\,\langle\alpha_j|\,, \tag{3}$$

where diagonal terms ρ_{ii} give the probability of state i and off diagonal terms ρ_{ij} and ρ_{ji} describe coherence between states i and j. A unitary operation $\hat{U}$ operates on the density matrix via

$$\rho \rightarrow U\rho U^{\dagger}, \tag{4}$$

where $U^{\dagger}$ is the Hermitian conjugate, and measurement of some observable Γ operates on ρ according to

$$\rho \rightarrow P_{\Gamma}[\rho] \tag{5}$$

where $P_{\Gamma}[\rho] = \sum_i |\gamma_i\rangle\,\langle\gamma_i|\,\rho\,|\gamma_i\rangle\,\langle\gamma_i|$ and $|\gamma_i\rangle$ are eigenvectors of Γ. For the sake of simplicity, unitary operators in the algorithm will be described in terms of operators acting on the ket side of the density matrix.

A common way to represent a quantum bit is to map the two state system to the surface of a Bloch sphere. Projections of the quantum bit onto some chosen axis, or rotations about that axis, can be made using Pauli spin operators

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{6}$$

$$\sigma_y = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}, \tag{7}$$

and

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{8}$$

Although a general quantum bit may be mapped to any vector on the Bloch sphere, all bits used in this paper lie in the xz plane of this sphere at all times, so that a single bit is rotated only about the y axis. A quantum bit is chosen to be true if it has positive projection on the $\hat{z}$ axis, and false if it has negative projection.

In the Deutsch model, quantum computation consists of an arbitrary unitary operation applied to a prepared initial state, so that

$$|\psi\rangle \to U|\psi\rangle, \tag{9}$$

and

$$\rho \to U\rho U^\dagger. \tag{10}$$

As most quantum algorithms involve interference between multiple pathways, it is usually assumed that the system initially occupies a pure state, so that

$$\rho = |\psi\rangle\langle\psi|. \tag{11}$$

In the context of NMR quantum computing [24], the initial pure state may be replaced by a pseudo pure state

$$\rho = (1-\epsilon)\mathbb{I}/N + \epsilon|\psi\rangle\langle\psi|, \tag{12}$$

so that

$$U\rho U^\dagger \to (1-\epsilon)\mathbb{I}/N + \epsilon U|\psi\rangle\langle\psi|U^\dagger, \tag{13}$$

where the totally mixed distribution $\rho_M = \mathbb{I}/N$ acts as a spectator to the computation acting on the pure state component.

3 Programming model for irreversible operations

The general approach of this paper is to incoherently transfer population from an arbitrary initial state to a state or set of states which satisfy some 3SAT problem of interest. This is accomplished by constructing decimation gates – essentially, error correction loops – which transfer population from states which fail individual clauses to states which satisfy those clauses. In pseudocode, the algorithm can be described as

for each clause in the problem **do**
  Encode the satisfaction of the clause in a temporary scratch bit
  if the clause is not satisfied **then**
    Use the scratch bit to perform a controlled rotation of the variables in the clause
  end if
  Clean up temporary scratch bits by measurement.
end for

Here the controlled rotation causes population loss for states which fail a particular clause. States which fail many clauses lose population more quickly than states which fail only a few, and solution states, which satisfy all clauses, act as population sinks.

The issue of reversibility arises due to the goal of mapping $N = 2^{N_V}$ initial states, where N_V is the number of variables, to a set of solution states which may be far smaller. Because a unitary operator maps a given initial state to a unique intermediate state at every point in the process,

$$\langle s'|U^\dagger U|s\rangle = \langle s'|s\rangle = \delta_{s,s'}, \tag{14}$$

nonunitary operations are necessary to perform a many to one mapping.

Such a mapping is achieved irreversibly through the introduction of temporary "scratch bits," which are introduced in prepared states before being entangled with the states of the

variable bits and finally deleted by nonselective measurement. The programming model which is used is to encode the outcome of a conditional statement in the projection of a scratch bit onto some axis (the "axis of truth"), then use the state of the scratch bit to perform controlled rotations of other bits.

Once the conditional rotations have been performed, the information necessary to reverse the process remains encoded in the projection of the scratch bit onto the axis of truth. Measuring the projection of the scratch bit along this axis would resolve which branch of the conditional statement was executed, and destroy the probability for states which correspond to the other branch. Instead, this information is deliberately destroyed by nonselectively measuring the scratch bit's projection onto an axis orthogonal to the axis of truth.

The nonunitarity of the measurement operator means that computations following this model will in general be nonunitary, and will yield mixed rather than pure quantum states. In contrast to the pseudo-pure states of NMR quantum computing, these mixed states have no special form, and the totally mixed density matrix $\rho_M = \mathbb{I}/N$ is not in general unchanged by the computation.

Choosing the axis of truth to be $\hat{z}$, logical operators are constructed so that the state of a scratch bit is $|+z\rangle$ if an expression is true and $|-z\rangle$ if the expression is false. Measuring the projection of the scratch bit along this axis would then resolve the truth of the expression, and thereby gather information about the state of the variables in question. In contrast, this information is lost if the projection of the scratch bit is measured along the $\hat{x}$ axis instead. Recalling that $|+z\rangle = (|+x\rangle + |-x\rangle)/\sqrt{2}$ and $|-z\rangle = (|+x\rangle - |-x\rangle)/\sqrt{2}$, measuring the state of the scratch bit along the $\hat{x}$ axis yields $|+x\rangle$ or $|-x\rangle$ with probability $1/2$, regardless of which branch was executed. Because either outcome is equally likely regardless of whether the expression is true or not, the measurement gives no information about the value of the expression or the variables in question. For a nonselective measurement, the measured value is not retained, so that the density matrix after measurement is the sum of the density matrices corresponding to each measurement outcome weighted by the probabilities of each outcome.

The goal of performing operations in this way is to construct a "decimation gate" corresponding to every three variable clause in the problem. The decimation gate incoherently transfers population from the set of states which fail to solve the clause to the set of states which do satisfy it. As indicated by the name, the goal is not to transfer all probability away from the states which fail a particular clause, but rather some fraction. In this way, states which fail many clauses lose population more quickly than states which fail only a few. Solution states, which satisfy all clauses, act as population sinks – probability which is transferred to these states is never lost due to the action of a decimation gate.

Expecting that nearly every state for the variable bits will fail at least one clause, the decimation gates are constructed in such a way that the satisfaction or failure of a particular clause is never measured. Instead, the decimation procedure is repeated until the density matrix approaches equilibrium. Provided that a solution exists, this occurs when all population has been transferred to the set of solution states. Once the total probability has been sufficiently concentrated in this set, the states of the variable bits can be measured, yielding a valid solution with high probability.

3.1 Comparison with fixed point quantum search

The algorithm described above has the property that population can be transferred to solution states, but never away. Because a solution state satisfies all clauses, it is a fixed point of the iteration. Conceptually, this is similar to the fixed point quantum search algorithm of Tulsi *et al* [25, 26, 27, 28, 29], in which a single oracle bit is used to rotate a source state $|s\rangle$ in a 2d subspace consisting of $|s\rangle$ and a solution state $|s^*\rangle$ which satisfies the problem. As described in [26], fixed point searches have the desirable properties that the initial state is guaranteed to evolve toward the target state, even when the algorithm is not run to completion, and any errors due to imperfect transformations in earlier iterations are wiped out by subsequent iterations, as long as the state remains in the problem defining space. However, existing fixed point searches have thus far offered no speed advantages relative to non fixed point algorithms. [26] yields running times comparable to classical search, while running times for [27] vary continuously from the $\mathbb{O}(\sqrt{n})$ running time characteristic of the Grover search algorithm, where n is the total number of states, to the $\mathbb{O}(n)$ of classical search with the variation of a damping parameter.

The present algorithm differs from existing fixed point algorithms by making use of information specific to 3SAT which is not available for the general case of an unordered database search – the condition that a valid solution must satisfy each clause in the problem. Rather than using a single oracle bit to encode the satisfaction of the entire problem, it makes use of multiple bits which encode the satisfaction of individual clauses. Rather than transferring population directly from a source state to a target state, a decimation operator acts to depopulate states which fail a particular clause, transferring population to other states which differ only by the variables contained in the clause. In this way, the flow of population mimics a breadth-first search for a state which satisfies the problem of interest. Ultimate transfer of population to the solution states occurs only indirectly, and may involve transfers to many intermediate states by means of many different pathways. The rates of depopulation due to failing individual clauses are user-controlled parameters with no *a priori* maximum value.

Each of these distinctions arises due to the additional problem specific information provided by 3SAT relative to a pure search of an unordered database, and collectively they account for the difference in convergence times. As shown in Section 7.1, the minimum fraction of population lost due to failure of an individual clause defines an important timescale for the decay of nonsolution population distributions. The condition that an eigendistribution decay more slowly than this uniquely defines a slowly decaying quasiequilibrium eigendistribution. In Section 7.2, varying the fractions of population lost due to failure of individual clauses projects the slowly decaying quasiequilibrium distribution onto quickly decaying eigendistributions, thereby increasing the rate at which population is transferred to the set of solution states. A detailed comparison with the Tulsi algorithm is given in Section 7.1.1.

4 Operators

4.1 Rotation operators

Because all qubits used in this paper lie in the xz plane at all times, the state of a particular bit can be represented by a vector on the unit circle. Quantum operators are written as the product of single bit rotation operators, which rotate a particular bit about the $\hat{y}$ axis, and controlled rotations, in which a target bit is rotated when a control bit is set to true

and left unchanged when the control bit is set to false. Measurement superoperators project a particular bit onto eigenstates with either positive or negative projection onto some axis. These three elementary operations are shown pictorially in Figure 1. Here the the single bit rotation operator is given by

$$R^n(\theta) = \exp[i\sigma_y^n\theta/2] = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \tag{15}$$

shown pictorially in Figure 1a)

The conditional rotation operator

$$\mathrm{CR}^{c\to t}(\theta) = R^t(\theta/2)\exp[i\sigma_z^c\sigma_y^t\theta/4] = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) & 0 & 0 \\ \sin(\theta/2) & \cos(\theta/2) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{16}$$

causes the state of target variable t to be rotated by angle θ when the projection of the control bit onto control axis $\hat{z}$ is positive and left unchanged if the projection is negative. These operations are shown pictorially in Fig. 1ci) and 1cii).

As the $\sigma_z^c\sigma_y^t$ operator of Eq. 16 may be difficult to realize physically, a controlled rotation may be more straightforward to implement using out of plane rotations. Defining a rotation about the x axis as

$$R_x(\theta) = e^{i\sigma_x\theta/2}, \tag{17}$$

Eq. 16 can also be written as

$$\mathrm{CR}^{c\to t}(\theta) = R^t(\theta/2)R_x^t(-\pi/2)e^{i\sigma_z^c\sigma_z^t\theta/4}R_x^t(\pi/2). \tag{18}$$

A second variation of the conditional rotation is to perform the rotation when the control bit is false and leave the target bit unrotated when the control bit is true. Defining

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{19}$$

it can be seen that

$$X^cCR^{c\to t}X^c = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\theta/2) & -\sin(\theta/2) \\ 0 & 0 & \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \tag{20}$$

4.2 Measurement Operators

The projection of bit n onto axis i is measured by projecting onto eigenstates $|\gamma\rangle$ of spin matrix

$$\sigma_i(\theta,\phi) = \cos(\theta)\sigma_z + \sin(\theta)\cos(\phi)\sigma_x + \sin(\theta)\sin(\phi)\sigma_y. \tag{21}$$

using projection superoperator

$$P_{\hat{\gamma}}^n[\rho] = \sum_i |\gamma_i\rangle\langle\gamma_i|\rho|\gamma_i\rangle\langle\gamma_i|, \tag{22}$$

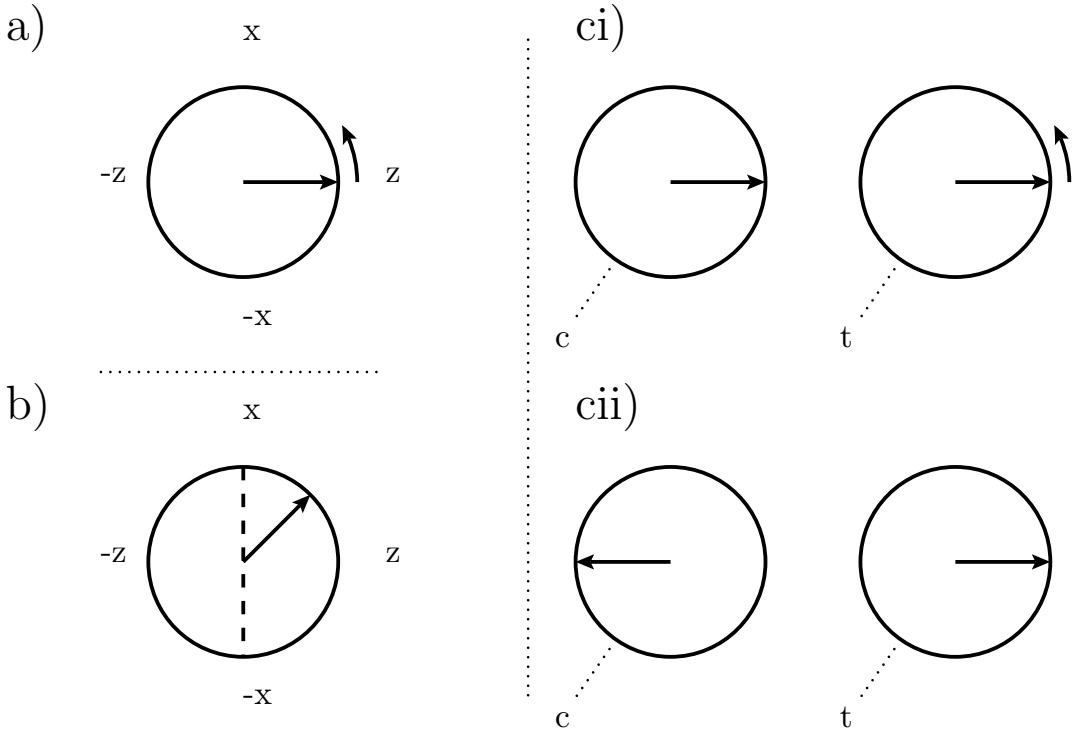

Fig. 1. Pictorial representation of elementary qubit operations used in this paper. a) Single bit rotation $R(\theta)$ rotates the bit by angle θ in the xz plane. b) Measurement operator (dotted line) nonselectively measures the projection of a bit onto a chosen axis. ci) Controlled rotation operator $CR^{c\to t}(\theta)$ rotates target bit t by angle θ when the control bit is true. cii) Controlled rotation operator $CR^{c\to t}(\theta)$ leaves target bit t unchanged when control bit c is false.

where γ_i are the eigenvalues of operator $\hat{\gamma}$.

For a nonselective measurement, in which the measured value of γ_i is not retained, Eq. 22 yields a new density matrix which sums over all measurement outcomes weighted by the probability of each outcome. This contrasts with a selective measurement, in which the measured value is retained and the summation consists of a single term.

For a scratch bit, where the outcome of a logical operation is encoded in the projection of the bit onto some axis of truth, the choice of measurement axis affects whether the outcome of the measurement yields any information regarding the states of the variables in the problem. If the axis of truth is chosen to be $\hat{z}$, with true represented by $|S\rangle = |+z\rangle$ and false by $|S\rangle = |-z\rangle$, then measuring along an axis at an angle θ relative to $\hat{z}$ in the xz plane yields eigenstates of the measurement operator

$$|+M\rangle = \cos(\theta/2)\,|+z\rangle + \sin(\theta/2)\,|-z\rangle \tag{23}$$

$$|-M\rangle = -\sin(\theta/2)\,|+z\rangle + \cos(\theta/2)\,|-z\rangle\,. \tag{24}$$

Because the measurement operator does not commute with truth operator σ_z, the outcome of the measurement yields only partial information regarding the output of the logical operator. For the special case when $\theta = \pi/2$, corresponding to measurement along the $\hat{x}$ axis, the probability is 1/2 to yield $|+x\rangle$ or $|-x\rangle$, regardless of whether the logical operator yields true or false.

To maintain the traditional meaning of measurement as obtaining information regarding the state of the system, it is useful to introduce new terminology. Let a deletion operation consist of projecting some bit onto eigenstates of a deletion operator (ie, measurement of a scratch bit along some axis). Deletion with measurement occurs when the detected state of the scratch bit uniquely determines the output of the logical operation – this occurs when the measurement axis is equal to the axis of truth. Deletion without measurement occurs when the detected state of the scratch bit yields no information about the state of the system. Choosing a measurement axis which is neither parallel nor perpendicular to the axis of truth yields deletion with partial measurement, so that the state of the scratch bit yields partial but not complete information about the output of the logical operation encoded in the scratch bit. Nonselective deletion occurs when the state of a bit is measured along some axis, but the value which is measured is not retained.

Using this terminology, the programming model described in section 3 can be described as insertion of a new bit in a prepared initial state, followed by entangling the state of the scratch bit with the states of variable bits, followed by nonselectively deleting the scratch bit without measurement. In this way, information regarding the initial state of the variable bits is lost, allowing for a transfer of population from many initial states to a few final states.

4.3 Logical operators

Conditional operations on a density matrix can be accomplished by introducing temporary scratch bits which store the outcome of logical operations. These bits can then be used to control operations on other bits, before being deleted by measurement.

A quantum OR gate can be defined through use of two temporary scratch bits

$$\begin{aligned} QOR^{c_1,c_2\to s_2} =& P_x^{s_1} R^{s_2}(\pi) X^{s_1} \mathrm{CR}^{s_1\to s_2}(\pi/2) X^{s_1} \\ & \mathrm{CR}^{c_2\to s_2}(\pi/2) \mathrm{CR}^{c_1\to s_2}(\pi/2) \\ & \mathrm{CR}^{c_2\to s_1}(\pi) \mathrm{CR}^{c_1\to s_1}(\pi) \end{aligned} \tag{25}$$

with circuit diagram

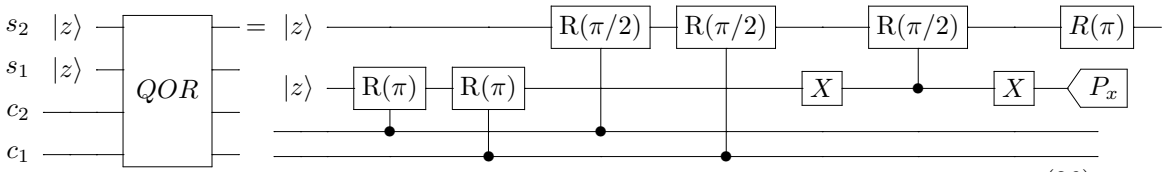

(26)

Here the second scratch bit, initialized to be true, remains true if either or both of the control bits are true, and is rotated to false otherwise. As before, the negation of a variable can be used to control the state of the scratch bit by an appropriate set of sign changes. Before its measurement, the projection of scratch bit s_1 onto the $\hat{z}$ axis encodes whether an even or odd number of input variables c_1 and c_2 are true. If only one is true, s_2 is rotated by an additional $\pi/2$, so that TT, TF, and FT yield the same projection for s_2 onto the $\hat{z}$ axis. Although not required for the current algorithm, an irreversible AND gate can be obtained by changing the angle of the $\mathrm{CR}^{s_1\to s_2}$ gate from $\pi/2$ to $-\pi/2$. Figures 2, 3, and 4 show the operation of the QOR gate pictorially as a series of rotations for initial states $|c_1c_2\rangle = |TT\rangle$, $|TF\rangle$, and $|FF\rangle$, and the final states are summarized in Table 5.

Use of two QOR gates and an intermediate scratch bit allows a second scratch bit to encode whether a particular quantum state satisfies three variable logical clause C

$$3\mathrm{OR}^{c_1,c_2,c_3\to s_4} = P_x^{s_1} \mathrm{QOR}^{c_3,s_2\to s_4} \mathrm{QOR}^{c_1,c_2\to s_2}, \tag{27}$$

with circuit diagram

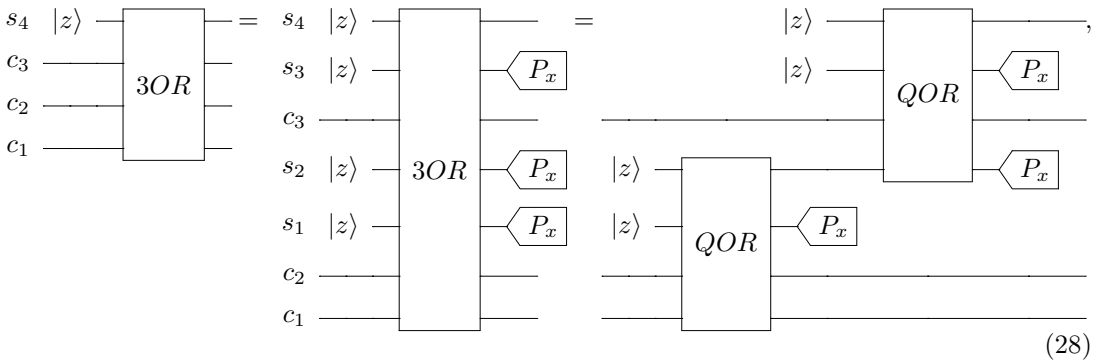

(28)

where c_1, c_2, and c_3 are the values of the variables in the clause or their negations. Initializing both scratch bits to be true, s_2 remains true if the clause is satisfied, and is rotated to false otherwise. Table 6 gives a truth table for the states of the scratch bits resulting from the QOR operator before they are measured.

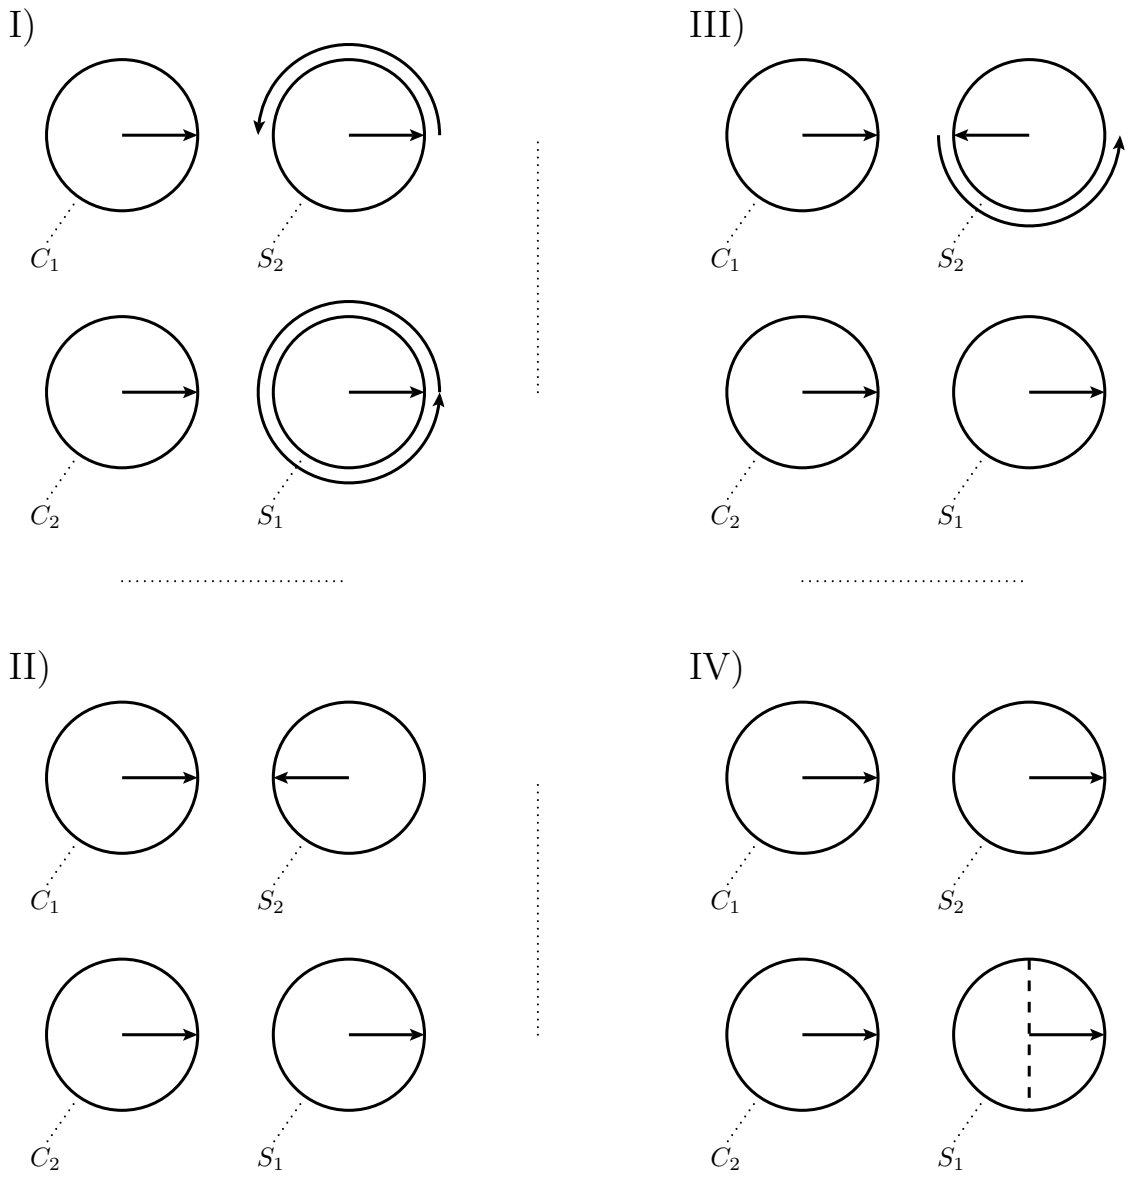

Fig. 2. QOR gate expressed as a series of bit rotations for initial state $|C_1\rangle = |T\rangle$, $|C_2\rangle = |T\rangle$. I) $\mathrm{CR}^{c_2 \to s_1}(\pi)\mathrm{CR}^{c_1 \to s_1}(\pi)$ yields 2π rotation of s_1, $\mathrm{CR}^{c_2 \to s_2}(\pi/2)\mathrm{CR}^{c_1 \to s_2}(\pi/2)$ yields π rotation of s_2. II) $X^{s_1}\mathrm{CR}^{s_1 \to s_2}(\pi/2)X^{s_1}$ yields no rotation of s_2. III) $R^{s_2}(\pi)$ rotates s_2 by π. IV) s_1 is deleted by nonselective measurement along the $\hat{x}$ axis.

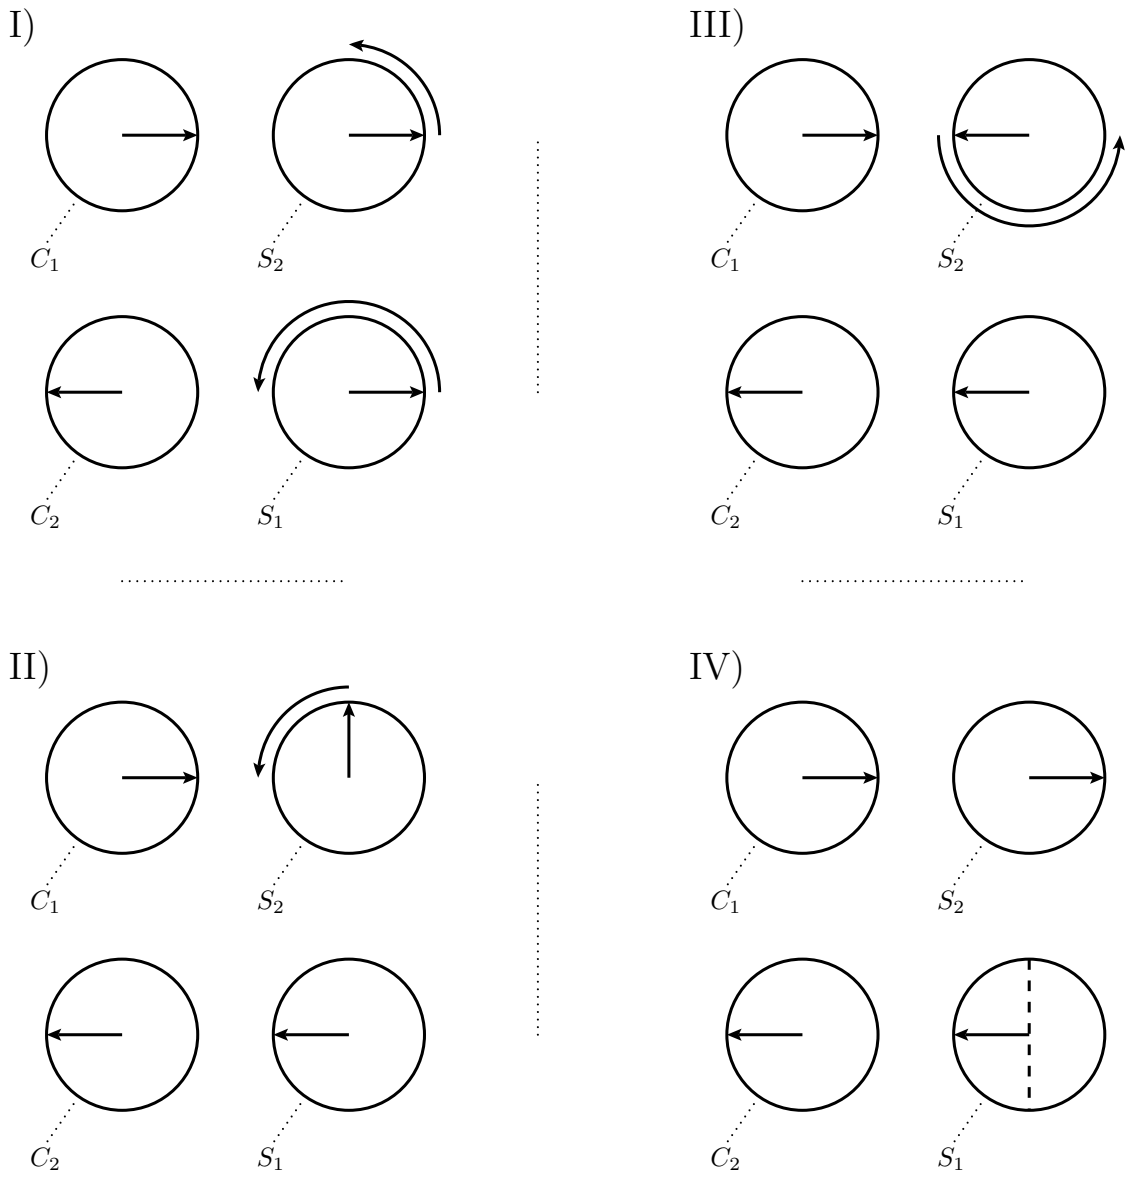

Fig. 3. QOR gate expressed as a series of bit rotations for initial state $|C_1\rangle = |T\rangle$, $|C_2\rangle = |F\rangle$. I) $\mathrm{CR}^{c_2 \to s_1}(\pi)\mathrm{CR}^{c_1 \to s_1}(\pi)$ yields π rotation of s_1, $\mathrm{CR}^{c_2 \to s_2}(\pi/2)\mathrm{CR}^{c_1 \to s_2}(\pi/2)$ yields $\pi/2$ rotation of s_2. II) $X^{s_1}\mathrm{CR}^{s_1 \to s_2}(\pi/2)X^{s_1}$ yields $\pi/2$ rotation of s_2. III) $R^{s_2}(\pi)$ rotates s_2 by π. IV) s_1 is deleted by nonselective measurement along the $\hat{x}$ axis.

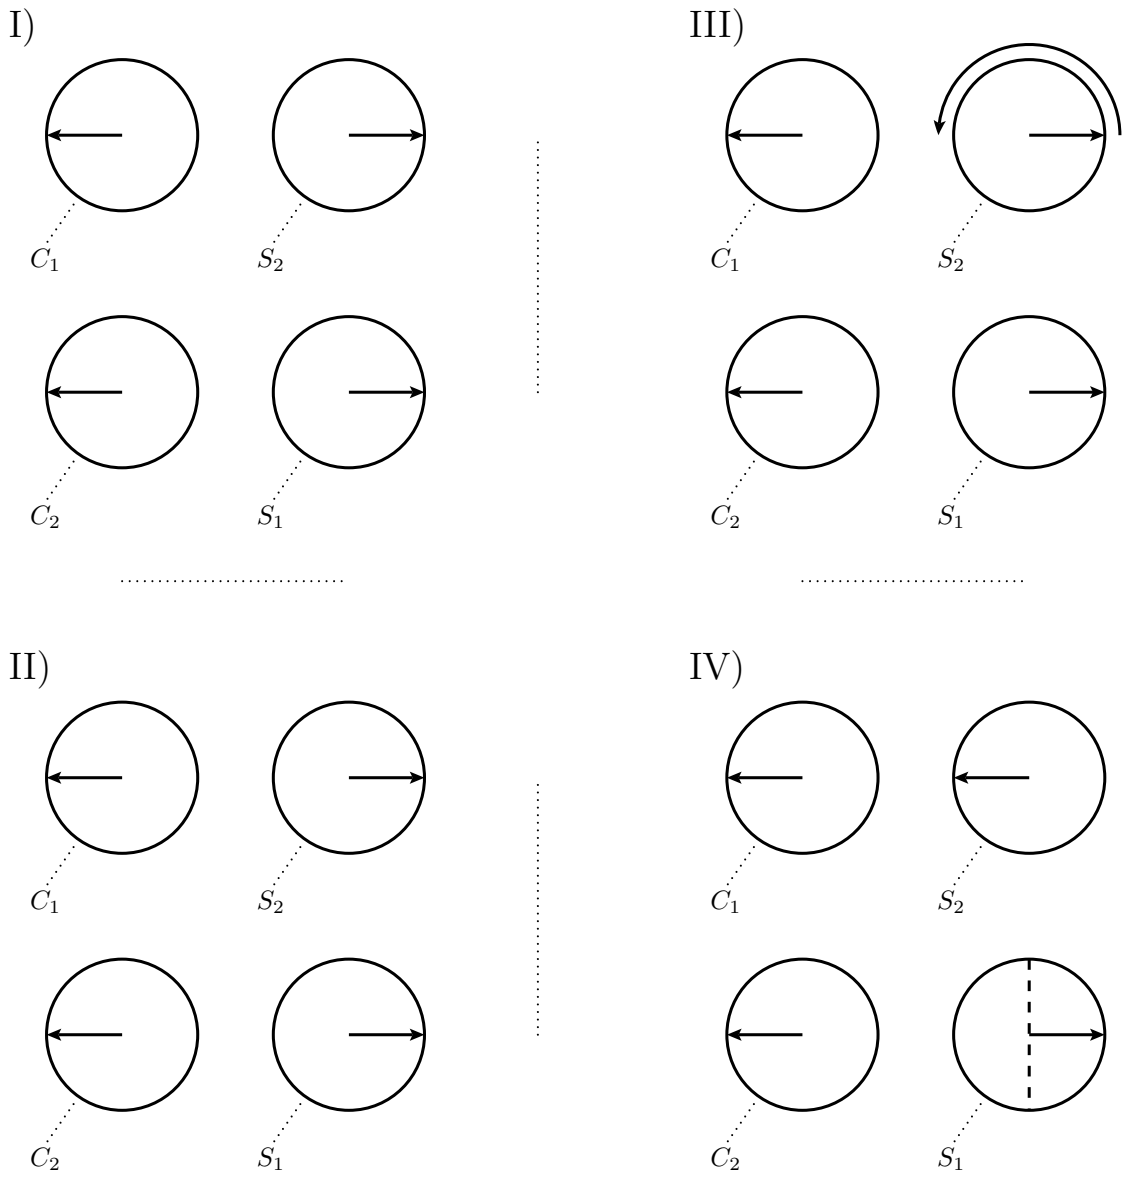

Fig. 4. QOR gate expressed as a series of bit rotations for initial state $|C_1\rangle = |F\rangle$, $|C_2\rangle = |F\rangle$. I) $\mathrm{CR}^{c_2 \to s_1}(\pi)\mathrm{CR}^{c_1 \to s_1}(\pi)$ yields no rotation of s_1, $\mathrm{CR}^{c_2 \to s_2}(\pi/2)\mathrm{CR}^{c_1 \to s_2}(\pi/2)$ yields no rotation of s_2. II) $X^{s_1}\mathrm{CR}^{s_1 \to s_2}(\pi/2)X^{s_1}$ yields no rotation of s_2. III) $R^{s_2}(\pi)$ rotates s_2 by π. IV) s_1 is deleted by nonselective measurement along the $\hat{x}$ axis.

Initial State	Final State
$\|c_1c_2;s_1s_2\rangle$	$\|c_1c_2;s_1s_2\rangle$
$\|TT;TT\rangle$	$\|TT;TT\rangle$
$\|TF;TT\rangle$	$\|TF;FT\rangle$
$\|FT;TT\rangle$	$\|FT;FT\rangle$
$\|FF;TT\rangle$	$\|FF;TF\rangle$

Fig. 5. QOR gate encodes the output of $(c_1 \vee c_2)$ in the state of s_2. Scratch bits s_1 and s_2 are initialized to be true ($|+z\rangle$) before being entangled with the states of c_1 and c_2 by the QOR operator. Here $|T\rangle = |+z\rangle$ and $|F\rangle = |-z\rangle$.

5 Decimation operators

A 3SAT problem is solved only by a state which satisfies every clause in the problem. Because a 3OR gate encodes the satisfaction of a particular clause in the state of scratch bit s_2, it can be used to achieve a progressive decimation of the population for any state which fails to satisfy the clause.

$$\begin{aligned} DEC(\theta) =& P_x^s X^s \mathrm{CR}^{s\to c_1}(2\theta) \\ &\times \mathrm{CR}^{s\to c_2}(2\theta) \\ &\times \mathrm{CR}^{s\to c_3}(2\theta) X^s \\ &\times 3\mathrm{OR}^{c_1,c_2,c_3\to s} \end{aligned} \tag{29}$$

with circuit diagram

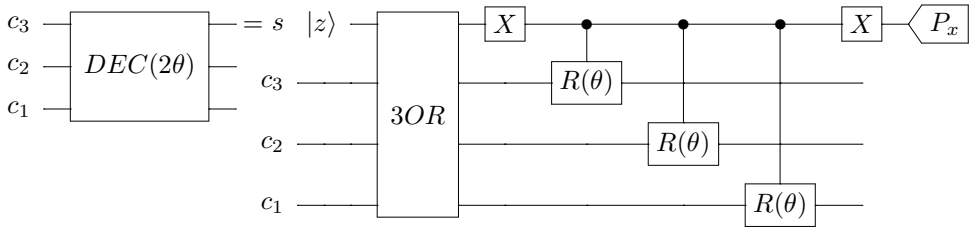

(30)

If the clause is satisfied, the states of variables X_{c_1}, X_{c_2}, and X_{c_3} are unaffected. If not, each variable is rotated by angle θ about the y axis. The decimation operator thus transfers population from any state which fails the clause to states which are identical to the original state except for the variables in clause C. Nonselectively measuring the scratch bit along the $\hat{x}$ axis makes this population transfer irreversible without giving any information regarding the satisfaction of the clause.

5.1 Deletion of scratch bits

The DEC operator combines a 3OR operator and a conditional rotation, so that the states of the variable bits are rotated if a given clause is unsatisfied, and left unchanged otherwise. The temporary scratch bits introduced by the QOR and 3OR operators are now deleted by nonselective measurement orthogonal to the axis of truth to yield a reduced density matrix for the variable bits alone.

Tables 5 and 6 show the states of the scratch bits after applying the QOR and 3OR operators. Because s_4 is used to perform a conditional rotation of the variable bits, $|FFF\rangle$

Initial State	Final State
$\|c_1c_2c_3; s_1s_2s_3s_4\rangle$	$\|c_1c_2c_3; s_1s_2s_3s_4\rangle$
$\|TTT; TTTT\rangle$	$\|TTT; TTTT\rangle$
$\|TTF; TTTT\rangle$	$\|TTF; TTFT\rangle$
$\|TFT; TTTT\rangle$	$\|TFT; FTTT\rangle$
$\|TFF; TTTT\rangle$	$\|TFF; FTFT\rangle$
$\|FTT; TTTT\rangle$	$\|FTT; FTTT\rangle$
$\|FTF; TTTT\rangle$	$\|FTF; FTFT\rangle$
$\|FFT; TTTT\rangle$	$\|FFT; TFFT\rangle$
$\|FFF; TTTT\rangle$	$\|FFF; TFTF\rangle$

Fig. 6. 3OR gate encodes the output of $(c_1 \vee c_2 \vee c_3)$ in the state of s_4, where s_1 and s_2 are the scratch bits corresponding to the first QOR gate in Eq. 27 and s_3 and s_4 the scratch bits corresponding to the second QOR gate. Each scratch bit is initialized to be true ($|+z\rangle$) before being entangled with the state of the input bits. Here the states of all four scratch bits are given prior to measurement, where $|T\rangle = |+z\rangle$ and $|F\rangle = |-z\rangle$.

is the only state which is rotated. To leading order in θ,

$$|FFF; TFTF\rangle \rightarrow (1-\frac{3}{2}\theta^2)\,|FFF; TFTF\rangle + \theta\,|FFT; TFTF\rangle + \theta\,|FTF; TFTF\rangle + \theta\,|TFF; TFTF\rangle \tag{31}$$

Applying this transformation to both the bra and the ket sides of the density matrix yields a matrix which contains indices for scratch as well as variable bits. Each scratch bit is now nonselectively measured along the $\hat{x}$ axis and traced over to give a reduced density matrix which contains indices for the variable bits only.

When tracing over the states of the scratch bits after measurement, all terms in which the state of a particular scratch bit is different on the bra side and the ket side will sum to zero. Writing

$$P_x^{s_n}[\rho] = \mathrm{Tr}_{s_n}[|x_{s_n}\rangle\langle x_{s_n}|\,\rho\,|x_{s_n}\rangle\langle x_{s_n}| + |-x_{s_n}\rangle\langle -x_{s_n}|\,\rho\,|-x_{s_n}\rangle\langle -x_{s_n}|], \tag{32}$$

it can be seen that

$$P_x^{s_n}[|T_{s_n}\rangle\langle T_{s_n}|] = P_x^{s_n}[|F_{s_n}\rangle\langle F_{s_n}|] = 1, \tag{33}$$

and

$$P_x^{s_n}[|T_{s_n}\rangle\langle F_{s_n}|] = P_x^{s_n}[|F_{s_n}\rangle\langle T_{s_n}|] = 0. \tag{34}$$

In particular, measuring scratch bit s_4 eliminates cross terms

$$P_x^{s_4}[|FFT; TTTT\rangle\langle FFT; TFTF|] = 0, \tag{35}$$

$$P_x^{s_4}[|FTF; TTTT\rangle\langle FFT; TFTF|] = 0, \tag{36}$$

$$P_x^{s_4}[|FFT; TTTT\rangle\langle FFT; TFTF|] = 0, \tag{37}$$

and so on, where the ket side arises due to an initial state which satisfies the clause and the bra side from the rotation of initial state $|FFF\rangle$ due to the DEC operator, or vice versa. Because state $|FFF\rangle$ is the only state yielding $s_4 = F$, all cross terms arising due to the conditional rotations of the decimation gate give zero contribution to the reduced density matrix.

A simple example of the decimation procedure is illustrated by using it to flip the state of a single variable. Here state $|\alpha\rangle$ evaluates to true for a particular clause, while state $|\alpha'\rangle$ evaluate to false. $|\alpha\rangle$ and $|\alpha'\rangle$ are identical except for the state of variable m. Because the satisfaction of the clause is encoded in the state of the scratch variable, the density matrix after the 3OR operator but before the conditional rotation can be written

$$\rho = \begin{pmatrix} \rho_{\alpha\alpha} & \rho_{\alpha\alpha'} \\ \rho_{\alpha'\alpha} & \rho_{\alpha'\alpha'} \end{pmatrix} \rightarrow \begin{pmatrix} \rho_{\alpha\alpha} & 0 & 0 & \rho_{\alpha\alpha'} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \rho_{\alpha'\alpha} & 0 & 0 & \rho_{\alpha'\alpha'} \end{pmatrix}, \tag{38}$$

with state ordering $|T,\alpha\rangle$, $|T,\alpha'\rangle$, $|F,\alpha\rangle$, $|F,\alpha'\rangle$.

When the state of the scratch bit is used to rotate the state of the variable bit according to $X^s \mathrm{CR}^{s\rightarrow m}(\theta) X^s$, the density matrix becomes

$$\rho = \begin{pmatrix} \rho_{\alpha\alpha} & 0 & \rho_{\alpha\alpha'}\sin(\theta) & -\rho_{\alpha\alpha'}\cos(\theta) \\ 0 & 0 & 0 & 0 \\ \rho_{\alpha'\alpha}\sin(\theta) & 0 & \rho_{\alpha'\alpha'}\sin^2(\theta) & -\rho_{\alpha'\alpha'}\cos(\theta)\sin(\theta) \\ -\rho_{\alpha'\alpha}\cos(\theta) & 0 & -\rho_{\alpha'\alpha'}\cos(\theta)\sin(\theta) & \rho_{\alpha'\alpha'}\cos^2(\theta) \end{pmatrix}. \tag{39}$$

The transfer of population from α' to α is completed by operator P_x^s, measuring the state of the scratch bit along an axis orthogonal to $\hat{z}$, the "axis of truth," to yield

$$\rho = \begin{pmatrix} \frac{1}{2}\left(\rho_{\alpha'\alpha'}\sin^2(\theta) + \rho_{\alpha\alpha}\right) & 0 & \frac{1}{2}(\rho_{\alpha\alpha'} + \rho_{\alpha'\alpha})\sin(\theta) & 0 \\ 0 & \frac{1}{2}\rho_{\alpha'\alpha'}\cos^2(\theta) & 0 & 0 \\ \frac{1}{2}(\rho_{\alpha\alpha'} + \rho_{\alpha'\alpha})\sin(\theta) & 0 & \frac{1}{2}\left(\rho_{\alpha'\alpha'}\sin^2(\theta) + \rho_{\alpha\alpha}\right) & 0 \\ 0 & 0 & 0 & \frac{1}{2}\rho_{\alpha'\alpha'}\cos^2(\theta) \end{pmatrix}. \tag{40}$$

Finally, tracing over the state of the scratch bit yields the new reduced density matrix

$$\rho = \begin{pmatrix} \rho_{\alpha\alpha} + \rho_{\alpha'\alpha'}\sin^2(\theta) & 0 \\ 0 & \rho_{\alpha'\alpha'}\cos^2(\theta) \end{pmatrix}, \tag{41}$$

where the transfer of population depends on the initial population of α' and the rotation angle, but not on the off-diagonal coherence terms $\rho_{\alpha\alpha'}$ or $\rho_{\alpha'\alpha}$.

5.2 Global decimation operator SATDEC

This decimation procedure can now be applied to each clause in the problem to create a global decimation operator

$$SATDEC(\mathfrak{O}, \vec{\theta}) = \prod_C DEC^C(\theta_C), \tag{42}$$

where $\mathfrak{O}$ gives the order of the clauses and θ_C gives the rotation angle for clause C. By construction, a state which satisfies every clause is unchanged by SATDEC, while every nonsolution state is decimated by at least one clause.

The algorithm described in Section 3 can now be implemented by repeated application of SATDEC to an arbitrarily chosen initial state.

6 Comparison with Schöning Algorithm

The algorithm described above is similar in spirit to the probabilistic classical algorithm of Schöning [30]. In the Schöning algorithm, the value of each variable is initially chosen at random. The satisfaction of each clause is then tested, and the failure of any clause is addressed by randomly flipping the state of one of the variables in the clause. If no satisfying assignment is found after testing $3N_V$ clauses, the process is restarted using a new random initial state.

The distinction between the two algorithms arises from the quantum nature of the system being operated on. In the current algorithm, the system consists of an incoherent superposition of many different states, while a classical algorithm is constrained to occupy a single state at a time. Schöning's algorithm can be emulated by measuring the states of each variable bit along the axis of truth after every decimation operation, thereby forcing the system to occupy a single state. Because the current algorithm does not transfer all probability away from the initial state, the steps of decimation followed by measurement must be repeated until one of the variable bit changes.

Here the act of measurement fundamentally alters the behavior of the algorithm. Treating each failed clause as a branch point, the current algorithm transfers a fraction of the probability to occupy the initial state along every branch, while the Schöning algorithm transfers the entire probability along a single branch chosen at random. In the language of computer science, the current algorithm is executing a parallel breadth first search for a solution state, while the Schöning algorithm executes a randomized depth first search – a random walk.

The flow of population due to a single DEC gate in the current algorithm can be emulated classically by summing over all random walks and all possible initial states, weighted by the initial probability of occupying a state and the probability of taking a particular walk. However, this classical emulation algorithm is very inefficient. Whereas the quantum algorithm requires a constant number of operations to implement a DEC gate, the classical algorithm requires a summation over all 2^{N_V} initial states. Thus, the existence of a polynomial time quantum algorithm does not imply the existence of a polynomial time classical emulation algorithm.

7 Convergence to Equilibrium

Provided that a solution exists, the decimation algorithm will eventually transfer all nonsolution probability to the set of solution states. The timescale on which this occurs is determined by the eigenvalues of the transfer matrix.

The probability distribution $P_s = \rho_{ss}$, corresponding to the diagonal elements of the density matrix, evolves due to a single application of SATDEC according to

$$\Delta P_s = \sum_{s'} T_{s' \to s} P_{s'} - T_{s \to s'} P_s, \tag{43}$$

where the transfer matrix T depends parametrically on the order of the clauses and the associated rotation angles: $T = T(\mathfrak{O}, \vec{\theta})$. Because the transfer of population away from a decimated state due a DEC gate goes as $\sin^2(\theta)$, to leading order it is quadratic in the decimation angle. Note that because $T_{s' \to s}$ is nonsymmetric, its eigenvectors are not orthogonal.

Because no solution state s^* is decimated, $P_{s^*} = 1$ for any solution state s^* is an eigenstate

of the transfer matrix with eigenvalue $\lambda = 1$. Let $\nu = 1 - |\lambda|$ be the decay rate associated with a nonsolution eigendistribution $\vec{\lambda}$, where $\vec{\lambda}[s]$ is the probability of state s for eigendistribution $\vec{\lambda}$.

Because total probability is conserved, the trace

$$\mathrm{Tr}(\vec{\lambda}) = \sum_s \vec{\lambda}[s] \tag{44}$$

for any eigendistribution with $\lambda \neq 1$ must be 0, while the condition that $P_s \leq 1$ at all times means that $|\lambda| \leq 1$.

7.1 The quasiequilibrium eigendistribution

The minimum decimation angle $\theta_0 \leq \theta_C$, where $\theta_C = \vec{\theta}[C]$ is the decimation angle corresponding to a particular clause, establishes a characteristic rate of decay for every nonsolution eigendistribution save one. Because $T_{s\to s'}$ is quadratic in the decimation angle, θ_0^2 gives the minimum fraction of a state's population which is lost due to being decimated, so that

$$\sum_{s'} T_{s\to s'} \geq \theta_0^2 \tag{45}$$

for all nonsolution states s.

The condition that $\nu << \theta_0^2$ is sufficient to uniquely determine a slowly decaying eigendistribution. Because $T_{s\to s'}$ is a nonnegative matrix, this is a special case of the Perron-Frobenius theorem. Writing Eq. 43 as an eigenvalue equation

$$\sum_{s'} T_{s'\to s} P_{s'} = T_{s\to s'} P_s - \nu P_s, \tag{46}$$

if $\nu << \theta_0^2$ it is negligible unless $T_{s\to s'} = 0$ for all s'; ie, unless s is a solution state.

Setting $\nu = 0$ in Eq. 46 and removing terms corresponding to solution states yields a system of equations with as many equations as unknowns.

$$\sum_{s'} T_{s'\to s} P_{s'} = \sum_{s'} T_{s\to s'} P_s. \tag{47}$$

If no solution states existed, this would be the equilibrium distribution. If solution states do exist, loss due to transfer of population from nonsolutions to solutions changes equilibrium to quasiequilibrium.

The quasiequilibrium eigenvector can be found in two steps. First, Eq 47 is solved for the nonsolution populations and normalized such that $\sum_s P_s = 1$. Populations for the solution states are found by setting

$$\Delta P_{s^*} = \sum_{s'} T_{s'\to s^*} P_{s'} \tag{48}$$

for all solution states s^*, and requiring $\Delta P_{s^*} = -\nu P_{s^*}$ for all s^* and $\sum_{s^*} P_{s^*} = -1$, so that the full eigendistribution is traceless.

Because the condition that $\nu << \theta_0^2$ uniquely determines the probability distribution for the slowly decaying eigendistribution, the contrapositive is also true: any nonsolution eigendistribution other than that determined by Eqs. 47 and 48 must have decay rates comparable to or larger than θ_0^2. Because these eigendistributions decay at a rate controlled by the user, they will be referred to as quickly decaying, in contrast with the slow decay of the quasiequilibrium distribution.

7.1.1 Comparison with the Tulsi algorithm

Applying the analysis of this section to the fixed point algorithm of Tulsi *et al* [26] sheds light on the difference between the two models, and offers a potential improvement in the Tulsi algorithm when applied to 3SAT. Whereas in this algorithm, the DEC operator depopulates states which fail a particular clause, with no consideration of whether the population is transferred to solution states or nonsolution states, the Tulsi algorithm employs the $\pi/3$ algorithm of Grover [25] to transfer population directly to a target state which is marked by an oracle bit. The minimum rate of transfer to a particular target state t is controlled by the minimum value of $||U_{ts}||^2$, where U is a unitary operator and U_{ts} the matrix element mapping source state s to target state t. Because

$$\sum_s ||U_{ts}||^2 = 1, \tag{49}$$

the minimum value is greatest when

$$||U_{ts}||^2 = 1/N, \tag{50}$$

where N is the total number of states. Thus, $T_{s\to s'} = 0$ when s and s' are both nonsolutions, and Eq. 45 becomes

$$\sum_{s^*} T_{s\to s^*} \approx M/N = f, \tag{51}$$

where M is the number of solutions and f the solution fraction.

Thus, the decay rate for eigenvectors in the Tulsi algorithm is limited by the solution fraction. In the worst case, $M = 1$ and $N = 2^{N_V}$, causing exponentially slow convergence. In contrast, in the current algorithm the decimation angles θ_C corresponding to particular clauses, which control the rate of decay for all eigendistributions save the quasiequilibrium distribution, are controlled by the user and do not change with the size of the problem. Although beyond the scope of this paper, a Tulsi algorithm which makes use of multiple oracle bits corresponding to the satisfaction of individual clauses rather than a single oracle bit corresponding to the satisfaction of the entire problem is worthy of investigation. The fraction of states satisfying an individual clause would then be $f_C = 7/8$, and would not change with the total number of states in the problem.

7.2 Back and Forth Iteration

Repeated application of SATDEC now causes the probability distribution to evolve on two timescales. The probability distribution at a given time can be written as a sum over eigendistributions

$$P_s = A_V \lambda_V^N \vec{V} + \sum_n A_n \lambda_n^N W_n + \sum_{s^*} A_{s^*} s^*, \tag{52}$$

where $\vec{V}$ is the quasiequilibrium eigendistribution defined by Eqs. 47 and 48 and $\vec{W}_n$ are the the remaining eigendistributions, with decay rates comparable to or greater than θ_0^2. Repeated application of the same transfer matrix results in a rapid "blowoff" in which nonsolution probability contained in quickly decaying eigendistributions is rapidly transferred to the set of solution states, followed by a long period of stagnation in which the nonsolution probability approaches the quasiequilibrium distribution, resulting in slow transfer to the solution states.

The slow decay of the quasiequilibrium distribution can be hastened by varying the decimation angles corresponding to different clauses. Because the slowly decaying eigendistribution is determined by $\vec{\theta}$ by way of Eq. 47, varying $\vec{\theta}$ non-adiabatically projects the original quasiequilibrium distribution into a new eigenbasis, in which the quasiequilibrium distribution may be substantially different. If $\vec{V}^{(1)}$ is the slowly decaying eigendistribution for transfer matrix $T^{(1)}$, and $\vec{V}^{(2)}$ the slowly and $\vec{W}_n^{(2)}$ the quickly decaying eigendistributions for transfer matrix $T^{(2)}$, where

$$T^{(2)}\vec{V}^{(1)} = (1 + A_V)\vec{V}^{(2)} + \sum_n A_n \vec{W}_n^{(2)}, \tag{53}$$

changing from transfer matrix $T^{(1)}$ to $T^{(2)}$ results in a second blowoff as the coefficients A_n decay to zero.

The coefficient A_V can be found by approximating the decay rate of the slowly decaying eigendistributions as zero, so that

$$T^{(1)}V^{(1)} = \quad V^{(1)} \tag{54}$$

$$T^{(2)}V^{(2)} = \quad V^{(2)}. \tag{55}$$

Let $\Delta T = T^{(2)} - T^{(1)}$ and $\Delta V = V^{(2)} - V^{(1)}$, so that Eq. 53 can be rewritten

$$A_V \vec{V}^{(2)} + \sum_n A_n \vec{W}_n^{(2)} = -T^{(2)}\Delta V. \tag{56}$$

Expanding the population difference

$$\Delta V = B_V V^{(2)} + \sum_n B_n \vec{W}_n^{(2)}, \tag{57}$$

in the eigenbasis of $T^{(2)}$, it follows that

$$A_V = -B_V = \frac{\Delta\vec{V} \cdot \vec{V}^{(2)}}{\vec{V}^{(2)} \cdot \vec{V}^{(2)}}, \tag{58}$$

so that the blowoff is largest when $\vec{V}^{(1)}$ and $\vec{V}^{(2)}$ are most dissimilar.

7.3 Back and Forth iteration

A simple procedure for transferring population away from the slowly decaying eigendistribution is to randomly vary the decimation angles corresponding to particular clauses for one iteration of SATDEC, then revert to the original decimation angles to recover the original transfer matrix for the next iteration. Let $T^{(1)} = T(\vec{\theta}^{(1)})$ be the transfer matrix corresponding to setting

$$\vec{\theta}^{(1)}[C] = \theta_0 \tag{59}$$

for every clause C, and $T^{(2)} = T(\vec{\theta}^{(2)})$, where

$$\vec{\theta}^{(2)}[C] = \theta_0(1 + \eta\sigma_c), \tag{60}$$

and σ_c is chosen to be 1 or -1 with equal probability.

Writing the transfer matrix $T = \prod_C T^{(C)}$ as the product of transfer matrices due to individual clauses C, note that $T^{(C)}$ is quadratic in θ_C,

$$T^{(C)} \propto \theta_C^2 \to \theta_0^2(1 + 2\eta\sigma_C). \tag{61}$$

The change in population for state s for the slowly decaying eigendistribution due to this change in the transfer matrix can be written

$$\Delta\vec{V}[s] = \sum_C \Delta V^{(C)}[s] + \Delta V^{(R)}[s], \tag{62}$$

where $\Delta V^{(C)}$ are first order corrections offsetting changes in $T^{(C)}$, while $\Delta V^{(R)}$ are residual corrections enforcing that the new quasiequilibrium distribution obeys Eq. 47.

To leading order, the change in population for state s due to a small change in the decimation angle is found by setting

$$\sum_{s'}(\vec{V}[s] + \Delta\vec{V}^{(C)}[s])(T^{(C)}_{s'\to s} + \Delta T^{(C)}_{s'\to s}) = \sum_{s'} \vec{V}[s]T^{(C)}_{s'\to s}, \tag{63}$$

so that after neglecting the doubly small term $\Delta T^{(C)}_{s'\to s}\Delta\vec{V}^{(C)}[s]$,

$$\frac{\Delta\vec{V}^{(C)}[s]}{\vec{V}[s]} = -\frac{\sum_{s'}}{\Delta T^{(C)}_{s'\to s}} T^{(C)}_{s'\to s} = -2\eta\sigma_C. \tag{64}$$

Subtracting these terms from Eq. 62 yields a linear system for the remaining corrections

$$\sum_{s',C}(\Delta\vec{V}^R[s]T^{(C)}_{s\to s'} - \Delta\vec{V}^R[s']T^C_{s'\to s}) = \sum_{s',C'\neq C}(\Delta\vec{V}^{(C')}[s']T^{(C)}_{s'\to s} - \Delta\vec{V}^{(C')}[s]T^{(C)}_{s\to s'}), \tag{65}$$

where the right hand side has expectation value 0.

The second step of the back and forth iteration changes all decimation angles back to their original values, so that

$$\Delta T^{(C)} \to -\Delta T^{(C)} \tag{66}$$

and

$$\Delta\vec{V} \to -\Delta\vec{V}, \tag{67}$$

and the slowly decaying eigendistribution $\vec{V}^{(2)}$ is projected into the eigenbasis of $T^{(1)}$.

Both changes to the transfer matrix result in a loss of total nonsolution probability, as population is transferred from the slowly decaying to the quickly decaying eigendistributions. The coefficient of the slowly decaying eigendistribution after the second transfer is given by

$$\frac{A'_V}{A_V} = (1 - \frac{\Delta\vec{V}\cdot\vec{V}^{(2)}}{\vec{V}^{(2)}\cdot\vec{V}^{(2)}})(1 + \frac{\Delta\vec{V}\cdot\vec{V}^{(2)}}{\vec{V}^{(2)}\cdot\vec{V}^{(2)}}). \tag{68}$$

Let

$$\alpha = \frac{\sum_{C,s}\Delta\vec{V}[s]\vec{V}[s]}{\sum_s \vec{V}[s]\vec{V}[s]} = \frac{\sum_C \sigma_C\eta\sum_s \vec{V}[s]\vec{V}[s]}{\sum_s \vec{V}[s]\vec{V}[s]} \tag{69}$$

and

$$\beta = \frac{\sum_s \Delta\vec{V}^R[s]\vec{V}[s]}{\sum_s \vec{V}[s]\vec{V}[s]}, \tag{70}$$

so that

$$\frac{A'_V}{A_V} = 1 - \alpha^2 - \alpha\beta - \beta^2. \tag{71}$$

Then

$$\alpha^2 = \frac{\sum_{C,C'} \sigma_C \sigma_{C'} \eta^2 \sum_{s,s'} \vec{V}[s]\vec{V}[s']}{\sum_{s,s'} \vec{V}[s]\vec{V}[s']}. \tag{72}$$

Because σ_C and $\sigma_{C'}$ are independent random variables with mean value 0,

$$\sum_C \sigma_C = 0 \tag{73}$$

and

$$\sum_{C,C'} \sigma_C^s \sigma_{C'}^s = \sum_C \sigma_C^2 = N_C^s \tag{74}$$

where N_C^s is the number of clauses failed by state s, so that $N_C^s \geq 1$ for every nonsolution state. From this it follows that the expectation value

$$\left\langle \frac{A'_V}{A_V} \right\rangle \leq (1 - \eta^2), \tag{75}$$

so that the loss of population from the slowly decaying eigendistributions due to a back and forth iteration cycle is controlled by a the user controlled parameter η.

7.4 Running time for the full algorithm

By transferring population from the slowly decaying quasiequilibrium distribution to the remaining, quickly decaying eigendistributions, back and forth iteration acts as a second efficient channel for population transfer to solution states, so that the overall rate of decay for nonsolution probability is determined by the user-controlled parameters θ_0 and η. Choosing $\eta = \sqrt{2}\theta_0$, the effective loss rate from the quasiequilibrium distribution is θ_0^2, the characteristic rate for the quickly decaying eigendistributions.

Assuming an initial nonsolution probability of 1, the number of SATDEC iterations required to yield solution probability P is a constant

$$N_{\text{iterations}} = \frac{\log(1-P)}{\log(1-\theta_0^2)}. \tag{76}$$

As each iteration of SATDEC requires a number of operations proportional to the number of clauses, the running time for the algorithm as a whole scales linearly with the size of the problem.

8 Summary and Conclusions

This paper has presented a new linear time quantum algorithm for the solution of the NP complete problem 3SAT. A longstanding goal of theoretical computer science, an efficient algorithm for 3SAT doubles as an efficient algorithm for all of class NP, including thousands of problems which are currently intractable using classical computers.

With many similarities to error correction algorithms, the present algorithm is well suited to implementation on fault-prone quantum processors. Off diagonal terms in the reduced density matrix describing the states of the variable bits play no role in the transfer of population to the set of solution states, removing the requirement for long lived coherence between different states of the system. Backwards transfer of population away from solution states and toward nonsolution states, which may arise due to imperfect execution of the algorithm or due to undesired physical processes in the processor, can be overcome by increasing the number of iterations, with no need to diagnose an error syndrome.

Aside from its practical significance, the relationship between algorithmic complexity and physical reversibility is worthy of further investigation. The scratch bits used in this algorithm are introduced in initially pure states before being entangled with the states of the variable bits and finally measured in such a way that two equally probable outcomes are possible. Thus, any decrease in entropy for the variable bits is accompanied by an increase in entropy elsewhere. Such a transfer of entropy is not possible using reversible operations on the variable bits alone. A deeper understanding of the connections between transfer of entropy and problems involving an exponentially large search space may help to shed light on still unresolved questions, such as the relationship between P and NP.

9 Appendix: A simple example of back and forth iteration

A simple example of back and forth iteration is obtained by solving a three variable 3SAT problem in which $|TTT\rangle$ is the only solution. This corresponds to performing DEC operations for clauses 0-6 of Table 7. Each state other than $|TTT\rangle$ fails a single clause, and to leading order in θ_C it is decimated by transferring population to the three states in which a single variable has been flipped. Choosing $\theta_C = \theta_0 = 0.1$ for all clauses and including terms to leading order in θ yields transfer matrix

$$T^{(1)} = \begin{array}{cccccccc} 0.97 & 0.01 & 0.01 & 0.0 & 0.01 & 0.0 & 0.0 & 0.0 \\ 0.01 & 0.97 & 0.0 & 0.01 & 0.0 & 0.01 & 0.0 & 0.0 \\ 0.01 & 0.0 & 0.97 & 0.01 & 0.0 & 0.0 & 0.01 & 0.0 \\ 0.0 & 0.01 & 0.01 & 0.97 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.01 & 0.0 & 0.0 & 0.0 & 0.97 & 0.01 & 0.01 & 0.0 \\ 0.0 & 0.01 & 0.0 & 0.0 & 0.01 & 0.97 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.01 & 0.0 & 0.01 & 0.0 & 0.97 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.01 & 0.0 & 0.01 & 0.01 & 1.0 \end{array} \tag{77}$$

where the state ordering is taken from table 7. The eigenvalues and eigendistributions of this transfer matrix are given by Table 8

#	Clause	State Decimated	States Populated
0	$(X_1 \vee X_2 \vee X_3)$	$\|FFF\rangle$	$\|FFT\rangle$, $\|FTF\rangle$, $\|TFF\rangle$
1	$(X_1 \vee X_2 \vee \neg X_3)$	$\|FFT\rangle$	$\|FFF\rangle$, $\|FTT\rangle$, $\|TFT\rangle$
2	$(X_1 \vee \neg X_2 \vee X_3)$	$\|FTF\rangle$	$\|FTT\rangle$, $\|FFF\rangle$, $\|TTF\rangle$
3	$(\neg X_1 \vee \neg X_2 \vee \neg X_3)$	$\|FTT\rangle$	$\|FTF\rangle$, $\|FFT\rangle$, $\|TTT\rangle$
4	$(\neg X_1 \vee X_2 \vee X_3)$	$\|TFF\rangle$	$\|TFT\rangle$, $\|TTF\rangle$, $\|FFF\rangle$
5	$(\neg X_1 \vee X_2 \vee \neg X_3)$	$\|TFT\rangle$	$\|TFF\rangle$, $\|TTT\rangle$, $\|FFT\rangle$
6	$(\neg X_1 \vee \neg X_2 \vee X_3)$	$\|TTF\rangle$	$\|TTT\rangle$, $\|TFF\rangle$, $\|FTF\rangle$
7	$(\neg X_1 \vee \neg X_2 \vee \neg X_3)$	$\|TTT\rangle$	$\|TTF\rangle$, $\|TFT\rangle$, $\|FTT\rangle$

Fig. 7. Each of the eight three variable states fails a single clause. To leading order, the DEC operator for this clause depopulates the state which fails the clause and populates the three states which differ from this state by a single variable.

Eigendistribution	λ	$\vec{\lambda}$
$\vec{s^*}$	1.0	$(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 1.0)$
$\vec{V}^{(1)}$	0.996	$(-0.165, -0.146, -0.146, -0.11, -0.146, -0.11, -0.11, 0.934)$
$\vec{W}_1^{(1)}$	0.98	$(-0.0, 0.577, -0.289, 0.289, -0.289, 0.289, -0.577, 0.0)$
$\vec{W}_2^{(1)}$	0.98	$(0.0, -0.086, -0.452, -0.537, 0.537, 0.452, 0.086, 0.0)$
$\vec{W}_3^{(1)}$	0.97	$(0.707, 0.0, -0.0, -0.354, -0.0, -0.354, -0.354, 0.354)$
$\vec{W}_4^{(1)}$	0.96	$(-0.0, -0.577, 0.289, 0.289, 0.289, 0.289, -0.577, 0.0)$
$\vec{W}_5^{(1)}$	0.96	$(0.0, 0.128, -0.551, 0.424, 0.424, -0.551, 0.128, 0.0)$
$\vec{W}_6^{(1)}$	0.944	$(-0.457, 0.403, 0.403, -0.305, 0.403, -0.305, -0.305, 0.162)$

Fig. 8. Eigenvalues and eigendistributions of $T^{(1)}$.

Setting $\eta = 0.2$ and $\vec{\sigma} = (1, 1, 1, 1, -1, -1, -1)$ changes the transfer matrix to

$$T^{(2)} = \begin{matrix} 0.957 & 0.014 & 0.014 & 0.0 & 0.006 & 0.0 & 0.0 & 0.0 \\ 0.014 & 0.957 & 0.0 & 0.014 & 0.0 & 0.006 & 0.0 & 0.0 \\ 0.014 & 0.0 & 0.957 & 0.014 & 0.0 & 0.0 & 0.006 & 0.0 \\ 0.0 & 0.014 & 0.014 & 0.957 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.014 & 0.0 & 0.0 & 0.0 & 0.981 & 0.006 & 0.006 & 0.0 \\ 0.0 & 0.014 & 0.0 & 0.0 & 0.006 & 0.981 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.014 & 0.0 & 0.006 & 0.0 & 0.981 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.014 & 0.0 & 0.006 & 0.006 & 1.0 \end{matrix}, \tag{78}$$

with eigenvalues and eigendistributions given by Table 9.

Note that in both Table 8 and Table 9, the decay rate for $\vec{V}^{(1)}$ and $V^{(2)}$ is small relative to $\theta_0^2 = 0.01$, while the decay rates for eigendistributions $\vec{W}_n$ are comparable to this value or larger.

Changing the transfer matrix from $T^{(1)}$ to $T^{(2)}$ projects $V^{(1)}$ into the new eigenbasis

$$\vec{V^{(1)}} = 1.005\vec{V}^{(2)} + 0.163\vec{W}_2^{(2)} - 0.085\vec{W}_3^{(2)} - 0.028\vec{W}_4^{(2)} - .005\vec{W}_6^{(2)}. \tag{79}$$

Once the coefficients of the quickly decaying eigendistributions $\vec{W}_n$ have decayed to zero, returning to the original transfer matrix projects $V^{(2)}$ into the original eigenbasis

$$\vec{V}^{(2)} = 0.956\vec{V}^{(1)} + 0.055\vec{W}_1^{(1)} - 0.13\vec{W}_2^{(1)} + 0.085\vec{W}_3^{(1)} - 0.014\vec{W}_4^{(1)} - 0.018\vec{W}_5^{(1)} 0.007\vec{W}_6^{(1)}, \tag{80}$$

so that the coefficient of $\vec{V}^{(1)}$ after the back and forth iteration is $1.005 * .956 = 0.96 \approx 1 - \eta^2$.

Eigenstate	λ	$\vec{\lambda}$
$\vec{s^*}$	1.0	$(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 1.0)$
$\vec{V}^{(2)}$	0.997	$(-0.101, -0.088, -0.088, -0.064, -0.234, -0.175, -0.175, 0.924)$
$\vec{W}_1^{(2)}$	0.984	$(0.0, -0.161, 0.161, -0.0, -0.0, -0.689, 0.689, 0.0)$
$\vec{W}_2^{(2)}$	0.98	$(-0.205, -0.322, -0.322, -0.399, 0.702, 0.19, 0.19, 0.166)$
$\vec{W}_3^{(2)}$	0.974	$(0.233, 0.02, 0.02, 0.034, 0.534, -0.546, -0.546, 0.25)$
$\vec{W}_4^{(2)}$	0.955	$(0.617, 0.043, 0.043, -0.654, -0.377, 0.069, 0.069, 0.189)$
$\vec{W}_5^{(2)}$	0.953	$(0.0, -0.626, 0.626, -0.0, -0.0, 0.329, -0.329, 0.0)$
$\vec{W}_6^{(2)}$	0.926	$(-0.492, 0.482, 0.482, -0.457, 0.165, -0.147, -0.147, 0.115)$

Fig. 9. Eigenvalues and eigendistributions of $T^{(2)}$.

10 Statement of Financial Interest

A patent application involving this work has been filed with the US Patent Office.

1. L. Fortnow, Communications of the ACM **52**(9), 78 (2009)
2. S.A. Cook, in *Proceedings of the third annual ACM symposium on Theory of computing* (ACM, 1971), pp. 151–158
3. L.A. Levin, Problemy Peredachi Informatsii **9**(3), 115 (1973)
4. R.M. Karp, *Reducibility among combinatorial problems* (Springer, 1972)
5. M.R. Garey, D.S. Johnson, *Computers and intractability*, vol. 29 (wh freeman, 2002)
6. D. Deutsch, in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 400 (The Royal Society, 1985), vol. 400, pp. 97–117
7. D. Deutsch, A. Barenco, A. Ekert, in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 449 (The Royal Society, 1995), vol. 449, pp. 669–677
8. A. Barenco, C.H. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J.A. Smolin, H. Weinfurter, Physical Review A **52**(5), 3457 (1995)
9. D.P. DiVincenzo, Physical Review A **51**(2), 1015 (1995)
10. E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, arXiv preprint quant-ph/0001106 (2000)
11. D. Aharonov, W. Van Dam, J. Kempe, Z. Landau, S. Lloyd, O. Regev, SIAM review **50**(4), 755 (2008)
12. R.P. Feynman, International journal of theoretical physics **21**(6/7), 467 (1982)
13. I. Buluta, F. Nori, Science **326**(5949), 108 (2009)
14. P.W. Shor, in *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on* (IEEE, 1994), pp. 124–134
15. P.W. Shor, SIAM journal on computing **26**(5), 1484 (1997)
16. L.K. Grover, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (ACM, 1996), pp. 212–219
17. L.K. Grover, Physical review letters **79**(2), 325 (1997)
18. J. Smith, M. Mosca, Handbook of Natural Computing pp. 1451–1492 (2012)
19. S. Jordan, Online: http://math. nist. gov/quantum/zoo/[Accessed 27 June 2013] (2011)
20. C.H. Bennett, E. Bernstein, G. Brassard, U. Vazirani, SIAM journal on Computing **26**(5), 1510 (1997)
21. B. Altshuler, H. Krovi, J. Roland, Proceedings of the National Academy of Sciences **107**(28), 12446 (2010)
22. I. Hen, A. Young, Physical Review E **84**(6), 061152 (2011)
23. E. Farhi, D. Gosset, I. Hen, A. Sandvik, P. Shor, A. Young, F. Zamponi, Physical Review A **86**(5), 052334 (2012)
24. J. Jones, Scalable Quantum Computers: Paving the Way to Realization pp. 139–154 (2005)
25. L.K. Grover, Physical Review Letters **95**(15), 150501 (2005)

26. T. Tulsi, L.K. Grover, A. Patel, Quantum Info. Comput. **6**(6), 483 (2006)
27. A. Mizel, Physical review letters **102**(15), 150501 (2009)
28. L. Hao, D. Liu, G. Long, Science China Physics, Mechanics and Astronomy **53**(9), 1765 (2010)
29. L. Gui-Lu, L. Yang, Communications in Theoretical Physics **50**(6), 1303 (2008)
30. U. Schöning, in *Foundations of Computer Science, 1999. 40th Annual Symposium on* (IEEE, 1999), pp. 410–414